

EXPLOITING MAP PLANS AS RESOURCES FOR ACTION

David Payton

Artificial Intelligence Center
Hughes Research Laboratories
3011 Malibu Canyon Road
Malibu, CA 90265

Abstract

When plans are used as programs for controlling the action of autonomous or teleoperated robots, their abstract representation can easily obscure a great deal of the critical knowledge that originally led to the planned course of action. In this paper, we highlight an autonomous vehicle experiment which illustrates how the information barriers created by abstraction can result in undesirable action. We then show how the same task can be performed correctly using plans as a resource for action. As a result of this simple change in outlook, we become able to solve problems requiring opportunistic reaction to unexpected changes in the environment.

I. Introduction

The teleoperation of robotic vehicles in space will require significant autonomous capabilities within the robotic vehicles. With long delays between sending commands from Earth and receiving them in space, telerobotic vehicles must be sufficiently responsive to their environment so that human operators need not be involved with every detail of the robot's motion. Ultimately, robotic vehicles must have such a high degree of autonomy that they may be capable of maneuvering through difficult terrain entirely on their own accord, forming their own plans to achieve user-specified goals.

In the endeavor to develop intelligent autonomous robotic agents capable of interacting with a dynamic environment, there has been a growing awareness that traditional planning methods may not be compatible with the demands for real-time performance. Recent efforts to re-evaluate the relationship between plans and action have led to alternative viewpoints in which plans are not primarily responsible for controlling a robot's behavior. Work by Brooks, for example, is aimed at avoiding the use of plans altogether [Br]. In this approach, intelligent action is a manifestation of many simple processes operating concurrently and coordinated through the context of a complex environment. While there is no tangible representation for plans in such a system, plans are implicitly designed into the system through the pre-established interactions between behaviors. Similarly, Agre and Chapman have shown how a system that determines its actions through the constant evaluation of its current situation can perform complex tasks that might otherwise have been thought to require planning [AC1]. Despite their emphasis on the theme that action is obtained by always knowing what to do at any instant, Brooks, Agre, and Chapman do not discard the notion that look-ahead and anticipation of future events are desirable activities. While these activities are

normally associated with planning, there is a difference in how the resultant "plans" are represented and used in their systems.

Agre and Chapman, for example, draw a sharp distinction between the concept of plans as communication and the more traditional views of plans as programs [AC2]. The key difference lies in the idea that plans must be constructed as a *resource* to the autonomous agent, not as an explicit set of instructions to be followed [Su]. As a resource, plans must serve as sources of information and advice to agents that are already fairly competent at dealing with the immediate concerns of their environment. In this sense, plans are used optionally, and serve only to enhance system performance. This is a significant departure from the conventional view of plans which puts them in the role of specifying a distinct course of action to systems which are often incapable of doing anything without them.

The differences between these two perspectives on planning are clearly evidenced when information from a map must be used to help guide an autonomous vehicle that must also make extensive use of sensors for detailed maneuvering and obstacle avoidance. In a plan-driven system, map-based plans are typically constructed to describe the optimal path that must be followed in order to arrive at a specified goal location. However, since the vehicle will invariably stray from the ideal path as it avoids sensed obstacles, the plan must be expressed in an abstract form that allows for error. In contrast, when map-based plans are represented for use as resources for action, this abstraction is not necessary. Instead, it is possible to make direct use of all information within the state-space of the map. As a result, information of all possible alternatives may be retained, allowing for flexible opportunistic behavior.

Our own experience with the DARPA Autonomous Land Vehicle (ALV) has led to some valuable insights into some of these issues. In a series of experiments performed by members of the Hughes Artificial Intelligence Center in August and December of 1987, a number of successful tests of autonomous cross-country navigation were performed using a system with integrated map and sensor-based control [Da] [KPR]. Some of the difficulties encountered in these experiments have pointed out certain consequences of the inappropriate use of abstraction that can occur in plan-driven systems. In this paper, we highlight one of these experiments to illustrate how the information barriers created by abstraction can lead to undesirable action. We then show how the same task can be accomplished without abstraction using plans as a resource for action, and we discuss how this approach may be extended for more complex problems.

II. The Misuse of Abstraction

In one of the cross-country experiments performed with the ALV we witnessed a surprising example of how easily plans can be misinterpreted in a plan-driven system. In this experiment, a very simple abstraction of a map-based plan was used to provide guidance to sensor-based obstacle avoidance behaviors. As shown in Figure [map plan], the basic mission objective was for the vehicle to get from one location to another while maintaining radio contact at all times. The map-based planner generated an appropriate route plan and abstracted a sequence of intermediate sub-goals to represent the critical points along this path. A portion of this sequence is illustrated in Figure [map plan] as Goals 1, 2, and 3. Note that the route had to veer specifically around one side of a rock outcrop in order to avoid loss of radio contact. To accomplish the mission, the sensor-

based behaviors had primary control of the vehicle so that all obstacles could properly be avoided. The behavior decisions, however, were always biased in favor of selecting a direction toward the current map sub-goal whenever possible. As soon as the vehicle got within a specified radius of its current sub-goal, that goal would be discarded and the next sub-goal would be selected. On paper and in simulation, it seemed that this approach would be effective.

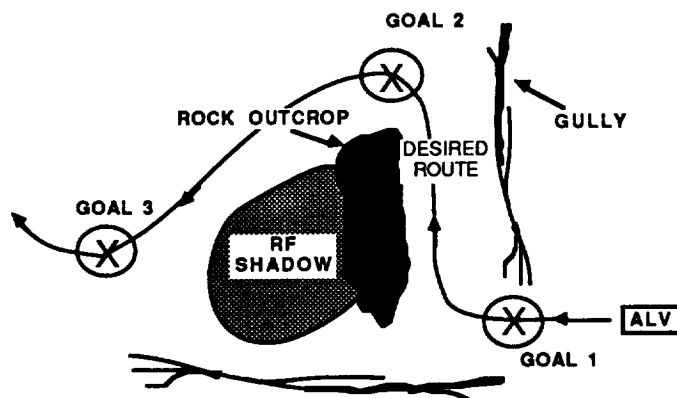


Figure [map plan]. An ALV route plan expressed as a sequence of intermediate goal points.

When we attempted to perform this mission with the ALV, the deficiencies of our method became strikingly clear. During the execution of this route, the vehicle achieved Goal 1 but then, because of local obstacles, was unable to turn appropriately to reach Goal 2. Figure [plan error] depicts the difference between the desired and actual routes. While this error is clearly apparent from the map data, the control behaviors had only the abstract route description as their guide, and this gave no indication that there was any problem with their action. Fortunately, contrary to our expectations, radio contact was not lost behind the obstacle. The mission could still be completed successfully if the vehicle were to move onward to Goal 3. Despite this new opportunity, however, the vehicle continued to persist toward Goal 2 because the abstract route description failed to give any indication that the original goal sequence was no longer suitable.

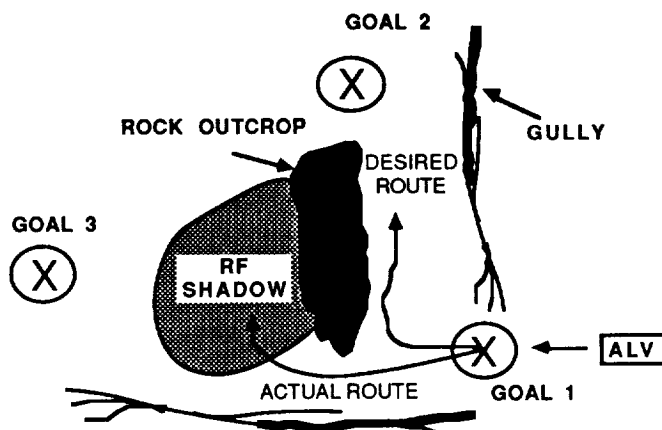


Figure [plan error]: Errant vehicle action while executing its route plan.

This example highlights the system's inability to take opportunistic advantage of unexpected situations when such situations are not properly accounted for in the abstract plan. We know from our understanding of the mission constraints that Goal 2 was merely an intermediate waypoint intended to keep the vehicle away from the RF shadow. Looking at the abstract plan in isolation, however, there is no way of knowing why a particular sub-goal has been established. The Goal 2 location could just as easily have been a critical choke point along the only path to Goal 3. It is only through our understanding of the underlying mission constraints that we can both identify the vehicle's failure to turn right and see the opportunity that arose as a result.

The apparent shortcoming of the abstract route plan is that it lacks environmental and mission constraints that are quite evident in the map. A more suitable plan would have explicated the concerns about staying out of the RF shadow. We therefore might wish to add more of this type of information to the plan. Once we start augmenting the plan, however, we have to ask how we might ever know when a sufficient amount of information has been added to prevent other types of mistakes. Consider, for example, the system's failure to realize that the intermediate sub-goal could be skipped when the opportunity arose. The problem arises because the true purpose of the sub-goal was never indicated. However, if the state-space of the plan could be expanded to include all the reasons for when and why the particular sub-goal was significant, then the location itself would become inconsequential. Consequently, the simple sequence of sub-goals is both an overspecification and an underspecification. The problem is inherent in any attempt to build an abstraction of the map data.

III. Avoiding Unnecessary Abstraction

In order to minimize the amount of information lost in forming a plan for action, it is best if all relevant knowledge is organized with respect to a given problem and then, without any further abstraction, provided in full for use in real-time decision-making. In order for this to be possible, the plan must no longer be viewed as a program for action, but rather, as a resource to help guide the decision-making process. When this viewpoint is adopted, there is no longer a need to translate plans into awkward representations for action. Instead, the original state-space in which the plan is formulated can be retained, enabling the plan to provide advice to decision-making processes whenever the current state of the system can be identified within that state-space. We refer to plans formulated and used in this manner as *internalized plans*, since they embody the complete search and look-ahead performed in planning, without providing an abstracted account of an explicit course of action [Pa].

The difference between the use of internalized plans and conventional abstracted plans is best illustrated in the context of the previous example. In contrast to the abstract route plan, consider a gradient description of a plan to achieve the same objectives. As illustrated in Figure [grad], there is no explicit plan shown, yet one can always find the best way to reach the goal simply by following the arrows. Such a representation would not ordinarily be thought to be a plan because it provides no specific course of action. As a resource for guiding action, however, the gradient field representation is extremely useful. No matter where the vehicle is located, and no matter how it strays from what might have been the ideal path, turn decisions can always be biased in favor of following the arrows.

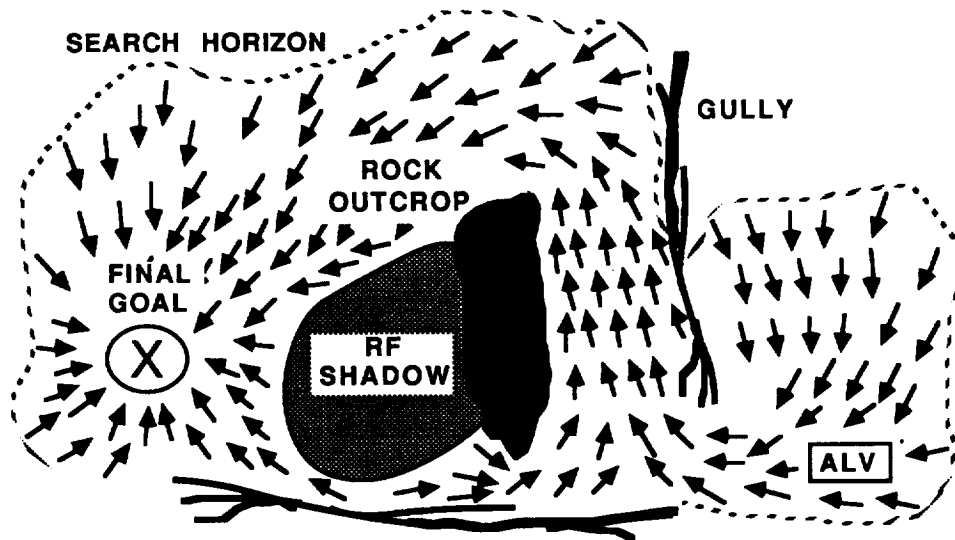


Figure [grad]: A gradient field representation provides one form of internalized plan.

Upon closer examination of Figure [grad], we can see not only how the mistake of entering the RF shadow could be avoided, but we see also how the system could be opportunistic should the vehicle happen to enter the shadow and be able to continue onward. First, when the vehicle had to make a choice between going left or right near the bottom of the rock outcrop, the gradient field would strongly bias its decision in favor of going right. If the vehicle got too close to the shadow on the left, the gradient field would actually be telling it to turn around. Further, should the vehicle happen to be forced to go below the rock outcrop and enter the RF shadow, then it would continue to be directed toward the final goal despite the radical deviation from its expected path. This type of behavior is opportunistic in that the vehicle is not constrained to reach any arbitrary pre-established sub-goals, and therefore all action can be directed exclusively toward achieving the mission objectives.

A more dramatic illustration of the difference between a conventional route plan and an internalized plan can be seen in problems requiring the attainment of any of several possible goals. This type of problem is often referred to as the "Post Office Problem" [Ed] because it can be likened to the task of finding the shortest route to the nearest of several post offices in a neighborhood. In the example shown in Figure [multi goal], the mission requires that the vehicle reach either of two distinct goal locations. The resultant gradient field is computed by propagating a search wavefront simultaneously from each of the two goals. As the wavefronts meet at a Voronoi edge, a ridge is created in the gradient field which will cause the vehicle to be guided toward one goal or the other depending on which side of the ridge it happens to be located.

Clearly, it would be difficult for an abstract route plan to capture the essence of choice contained in the gradient field representation. If we were to produce a route plan, we would invariably have to select a route to the closest goal, as shown in Figure [multi goal]. Once such a choice is made, however, we have discarded all that is known about the alternate goal even though that goal was nearly as close as the one selected. In contrast, by using the gradient field directly, the choice of

goals may be made during the execution of the mission. Without having made an *a priori* selection of goals, the best choice may be made at every instant in time, regardless of how the vehicle might stray while avoiding obstacles.

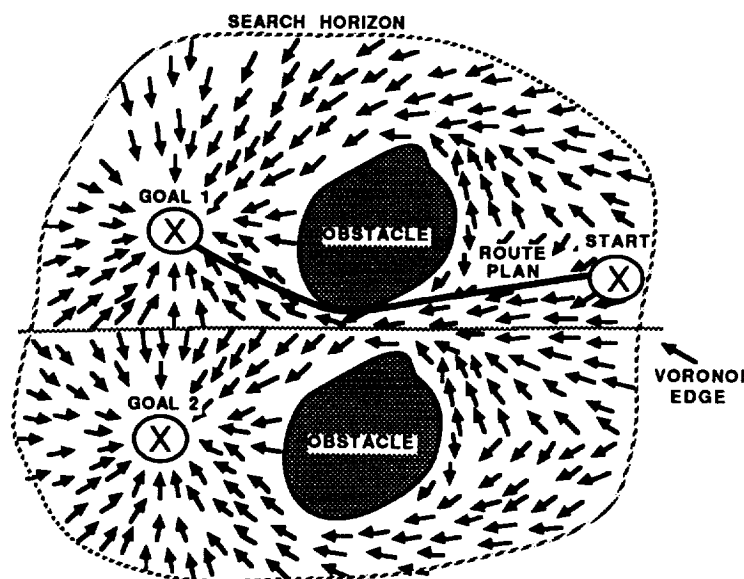


Figure [multi goal]: The gradient field provides a useful internalized plan for reaching either of two goals.

The gradient field is an ideal example of an internalized plan because the map-grid state-space in which the original problem is formulated is the same state-space in which the plan is represented. The gradient field, in fact, is a natural by-product of existing route planning algorithms [MPK]. These algorithms begin by assigning a cost to each grid cell of a digital terrain map. By associating high costs with locations that are undesirable according to mission criteria, a combination of mission constraints can be represented. Whether an A^* [Ni], or Dijkstra [Di] search algorithm is employed in the cost grid, the net result of the search is a score for each grid cell, indicating the minimum cost remaining to get from that cell to the goal. From any given grid cell, the best incremental step to get to the goal is the neighboring grid cell which has the lowest score. Ordinarily, when we use these scores to compute a standard route plan, we simply begin at the starting point and locally choose the lowest-score adjacent cell until we finally reach the goal. The record of our steps along the way gives us the minimum cost path to the goal. If we look at these scores in a slightly different way, we see that the best path to the goal from any grid cell may be determined by selecting the direction of the lowest-score adjacent cell. Thus, without any further abstraction, search in the map-grid can provide a useful resource for action.

IV. Using Plans as Resources

The method of use of a gradient field is an important factor in establishing it as an internalized plan representation. Since a digital terrain map generally cannot provide adequate resolution to

support detailed maneuvering around small obstacles, there is inevitably a need to incorporate the advice provided by the gradient field into real-time decision-making processes which are attending to immediate sensory data. While, ordinarily, a single abstract route plan is generated, some approaches have taken advantage of a gradient field in order to quickly generate new route plans should the constraints of an initial plan be violated [LMD] [CF]. Problems with establishing and monitoring these constraints, however, are still unavoidable. In contrast, use of the gradient field as an internalized plan requires that the real-time decision-making processes continuously attempt to locate the system within the state-space of the plan and bias each decision in favor of the recommended course of action. The absence of an explicit course of action means that no arbitrary plan constraints need be established or monitored. The plan is a resource, providing suggestions for preferred action but never actually controlling the system. If, for any reason, no suggestion is available from the plan, the real-time decision-making processes must proceed in a reasonable manner on their own accord.

Another vector field type of representation, the *artificial potential field*, appears superficially very similar to the gradient field and it also is used for robot navigation and obstacle avoidance [Kr][Kh][Ar]. The basic differences, though, between how these two types of representations are constructed and used sheds further light on what it means for a plan to serve as a resource for action. The computation of potential fields is generally based on a superposition model in which charges are distributed such that repulsive forces are generated near obstacles and attractive forces are generated near goals. Superposition allows the potential field vector at any point to be computed quickly by adding up the contributions from each charge. The resultant field, however, does not represent an optimal path, and may easily contain local minima and traps. In contrast, the gradient field is computed from a more time consuming graph search process. As a result of this search, the gradient field has no local minima and will always yield the set of all optimal paths to the goal.

A more significant distinction between gradient fields and potential fields, however, is in how they are used. Often, when potential field methods are employed for navigation, the potential field is used for direct control of action. All sensory information is compiled into a single representation which is suitable for modeling an appropriate distribution of charges. The local potential field forces are then continuously computed at the location of the vehicle, and these forces are used directly to compute the desired motion. On the other hand, as internalized plans, gradient fields are never used to provide direct control of the vehicle. Instead, they are merely an additional source of information provided to a set of real-time decision-making processes. Since these processes can make use of many disjointed representations of the world in order to control the vehicle, there is never a need for all features of the environment to be abstracted into a single representational framework.

It is helpful to view internalized plans as though they were sources of supplementary sensory input data. From this perspective, it is clear that action is not controlled by plans any more than it is by sensory input. Instead, the system must be viewed as an entity which interacts with its environment, responding to both internal and external information sources. The gradient field plan, for example, can be thought of as a phantom compass that always gives a general idea of the right way to go. Just like other sensors, data from this internal sensor influences action but is never used to the exclusion of other sensory data. At any given time, however, a single information source can have significant influence over system behavior if need be. Just as an external sensor can be used to ensure that the vehicle never runs into obstacles, an internalized plan can be used to ensure that

mission constraints are not violated. Thus, despite the fact that there is no top-down control, the system can adhere to high level mission requirements.

V. Multiple Internalized Plans

A significant advantage of using internalized plans as resources for action is that it is possible to use multiple internalized plans simultaneously. Each plan can contribute an additional piece of advice which can enhance the overall performance of the system. In this way, different plans may be formulated in incompatible state-spaces without the need to merge these state-spaces through abstraction.

We can consider as an example, the combined use of map-based plans with plans based on symbolic mission constraint data. In the case of the RF shadow problem, a constraint to maintain radio contact may be derived from mission knowledge. If this knowledge is used in conjunction with a signal strength sensor, then whenever the vehicle enters an RF shadow, it can immediately back up in order to regain contact. In the absence of such problems, the gradient field produced from map data can constantly provide advice on which way to go. An unexpected loss of radio contact would then be treated much like an encounter with an obstacle. The vehicle would have to make special maneuvers in order to regain contact and ensure that the same mistake would not be repeated. After this, the map-based plan would regain primary influence.

There are also many cases in which it might be desirable to use multiple internalized plans formulated within the same state-space. For example, a gradient field plan could be augmented with information about the amount of fuel and time required to get from each grid-square to the goal. While this information could not directly indicate a course of action, it might allow available fuel and time resources to be monitored constantly and compared with expected needs. If there were barely enough fuel to succeed but plenty of time available, the vehicle might be able to switch to a simple fuel conserving strategy such as reducing its speed. If time and fuel were both in short supply, the gradient field might need to be re-computed, placing more emphasis on conserving fuel and time resources and possibly less emphasis on other factors such as vehicle safety.

Another form of internalized plan exploits the map as a resource for action by probing it directly during execution. As the vehicle is traveling, the portion of the map corresponding to the area just in front of the vehicle is examined to determine what types of features should be detected. This understanding of the local environment can have a direct bearing on how sensor data is interpreted for action. Remember, for example, the problem illustrated earlier in Figure [plan error]. Here, one of the main reasons the vehicle failed to avoid the RF shadow was that its sensors indicated a clear path in this area. This error could be overcome by differentiating between obstacles that are observable and those that are not, and then appropriately discounting sensor readings that are known to be inapplicable. Thus, by treating the map as if it were a sensor, the value of real sensor data can be greatly enhanced.

A great diversity of behavior may also be gained by dynamically combining information from multiple gradient fields. Consider, for example, two independent gradient fields, one which can guide a vehicle along a safe, well hidden route, and another which can lead the vehicle to nearby observation points. We can imagine that the vehicle is guided by the safe gradient field until the time comes for it to make an observation. Then, the gradient field for getting to observation points would

become the primary guiding factor. Such a gradient field, formed similar to the field in Figure [multi goal], would lead the vehicle to the nearest of several possible observation points. Once an observation point had been reached and observation data collected, the safe gradient field would again be used for guidance. Using such a combination of internalized plans allows the performance of tasks that would be difficult to accomplish with a symbolic plan. Without an explicit plan for action, it is the interplay between the vehicle and its environment that determines how the mission will ultimately be carried out.

VI. Conclusion

Although abstraction is necessary if we are to provide organization and structure to the vast amounts of information available to an intelligent agent, we have seen examples in which the abstraction of plans can obscure their true intent and result in serious failures. In light of these issues we must ask whether forming the abstraction was really necessary or whether it was merely an artifact of an approach in which plans are regarded as programs rather than as resources for action. Using internalized plans, we have shown that with no abstraction of the map-based plan, we can obtain an ideal resource for action.

Just as the grid of a digital terrain map is an abstraction of the Earth's surface, abstraction may be used to create other state-spaces which are suitable to use for planning. In many cases, however, it may be best not to attempt the fusion of information from different sources if an excessive degree of abstraction is required to do so. Instead, state-spaces should be formed to suit the type of information available, and once planning is performed in these state-spaces, no further abstraction of the results should be performed. The unabstracted product of planning search provides a measure of desirability for transitions from one state to the next, and this measure may be used directly as a resource for action.

VII. Acknowledgments

Many of the concepts presented in this paper were fostered through frequent discussions with David M. Keirsey, J. Kenneth Rosenblatt, and Charles Dolan. I am extremely grateful for their contributions and insights. The advice of Joseph Mitchell and Rodney Brooks was also very helpful in bringing to light many important areas of related work. Many thanks also to my wife Karen and Jimmy Krozel for their editing assistance.

VIII. References

- [AC1] Agre, P., and D Chapman, "Pengi: An implementation of a Theory of Activity", *Proc. of the Sixth National Conf. on Artificial Intelligence*, Seattle, Washington, July, 1987, pp. 268-272.
- [AC2] Agre, P., and D. Chapman, "What are plans for?" AI Memo 1050, MIT Artificial Intelligence Laboratory, 1987.
- [Ar] Arkin, R., "Motor Schema Based Navigation for a Mobile Robot: An Approach to Programming by Behavior", *IEEE Conf. on Robotics and Automation*, March 1987, pp. 264-271.

- [Br] Brooks, R. A., "Intelligence Without Representation," *Preprints of the Workshop in Foundations of Artificial Intelligence*, Endicott House, Dedham, MA, June, 1987.
- [CF] Chan, Y.K., and M. Foddy, "Real Time Optimal Flight Path Generation by Storage of Massive Data Bases", IEEE National Aerospace and Electronics Conf. (NAECON), Dayton OH, May 1985.
- [Da] Daily, M., J. Harris, D. Keirsey, K. Olin, D. Payton, K.Reiser, J. Rosenblatt, D. Tseng, and V. Wong, "Autonomous Cross-Country Navigation with the ALV," *Proceedings of DARPA Knowledge-Based Planning Workshop*, Austin, Texas, December 1987, (also appearing in *Proceedings of IEEE Conference on Robotics and Automation*, Philadelphia, PA., April, 1988.)
- [Di] Dijkstra, E.W., "A Note on Two Problems in Connection with Graph Theory", *Numerische Mathematik*, Vol 1, 1959, pp. 269-271.
- [Ed] Edelsbrunner, H., *Algorithms in Combinatorial Geometry*. Springer-Verlag, Berlin, 1987, pp 298-299.
- [KPR] Keirsey, D.M., D.W. Payton, and J.K. Rosenblatt, "Autonomous Navigation in Cross Country Terrain", *Proceedings Image Understanding Workshop*, Boston, MA, April, 1988.
- [Kh] Khatib, O. "Real Time Obstacle Avoidance for Manipulators and Mobile Robots", *IEEE Conf. on Robotics and Automation*, March 1985, pp. 500-505.
- [Kr] Krogh, B.H., "A Generalized Potential Field Approach to Obstacle Avoidance Control", *Robotics International Robotics Research Conference*, Bethlehem, PA, August 1984.
- [LMD] Linden, T.A., J.P. Marsh, and D.L. Dove, "Architecture and Early Experience with Planning for the ALV", *IEEE International Conf. on Robotics and Automation*, April, 1986, pp. 2035-2042.
- [MPK] Mitchell, J.S.B., D.W. Payton, and D.M. Keirsey, "Planning and Reasoning for Autonomous Vehicle Control," *International Journal for Intelligent Systems*, Vol. 2, 1987.
- [Ni] Nilsson, N.J., *Problem Solving Methods in Artificial Intelligence*, McGraw-Hill, 1971.
- [Pa] Payton, D.W., "Internalized Plans: a representation for action resources", *Workshop on Representation and Learning in an Autonomous Agent*, Lagos, Portugal, Nov. 1988.
- [Su] Suchman, L., "Plans and Situated Actions: the problem of human machine communication", Cambridge University Press, 1987.